arXiv:1412.7437v1 [quant-ph] 23 Dec 2014

MIT–CTP 4619

Compressibility of positive semidefinite factorizations and quantum models

Cyril J. Stark and Aram W. Harrow
Center for Theoretical Physics, Massachusetts Institute of Technology,
77 Massachusetts Avenue, Cambridge MA 02139-4307, USA
(Dated: December 24, 2014)

We investigate compressibility of the dimension of positive semidefinite matrices while approximately preserving their pairwise inner products. This can either be regarded as compression of positive semidefinite factorizations of nonnegative matrices or (if the matrices are subject to additional normalization constraints) as compression of quantum models. We derive both lower and upper bounds on compressibility. Applications are broad and range from the statistical analysis of experimental data to bounding the one-way quantum communication complexity of Boolean functions.

I. INTRODUCTION

The following situation is ubiquitous in quantum information: a d-dimensional state ρ_x is prepared and a measurement E_y is performed with POVM elements $E_{y,z}$, with $x \in [X], y \in [Y], z \in [Z]$. The resulting conditional probability distribution of these outcomes $p(z|x,y) = \mathrm{tr}(\rho_x E_{y,z})$ can be expressed as a matrix $\mathcal{D}$ with $\mathcal{D}_{x;y,z} = p(z|x,y)$. The forward problem of computing $\mathcal{D}$ given $\{\rho_x\}, \{E_y\}$ is straightforward, but often we need to solve the *inverse* problem of finding states and measurements compatible with a given $\mathcal{D}$. This problem of finding *quantum models* is in general underdetermined, but we can constrain the problem by minimizing the dimension d of the model. Similarly, if given an accuracy parameter ε, we could ask for the minimum d for which a d-dimensional quantum model can approximate each entry of $\mathcal{D}$ to additive error ε.

Two important examples where this problem occurs are:

- *Inferring quantum models.* Suppose we perform an experiment where x and y are classical inputs (e.g. x selects the state and y the measurement) and z is classical observable. We will sample from the distributions $p(z|x,y)$ and thereby learn an approximation of $\mathcal{D}$. However, we do not have direct access to the underlying quantum states or measurements. In reconstructing these states and measurements, it is natural (e.g., to prevent overfitting) to posit the simplest model consistent with the data, which in the absence of other information would mean the quantum model of lowest dimension.

- *One-way quantum communication complexity.* Suppose that Alice and Bob would like to jointly compute a function $f : [X] \times [Y] \mapsto [Z]$ when Alice is given input x and Bob is given input y. The most general protocol consists of Alice sending ρ_x to Bob who performs POVM E_y and outputs the outcome z that he obtains. The log of the minimum dimension achieving $p(z|x,y) = \delta_{z,f(x,y)}$ (resp. $p(z|x,y) \approx \delta_{z,f(x,y)}$) is the exact (resp. approximate) one-way quantum communication complexity of f.

Quantum models for conditional probability distributions are an example of the more general idea of a psd (positive semidefinite) factorization. If M is a matrix with nonnegative entries, then a psd factorization of M is a collection of d-dimensional psd matrices $\{A_x\}, \{B_y\}$ such that $M_{xy} = \mathrm{tr}(A_x B_y)$. The psd-rank of M is the minimum d for which this is possible; approximate versions can also be defined [6, 8]. See [5] for a recent review of psd-rank. One application of psd-rank is to the case when M is the slack matrix of a polytope, in which case the psd-rank is the smallest possible dimension of a semidefinite program representing the polytope; here too approximate versions of this relation are known [8]. Our problem of finding low-dimensional quantum models is a special case of the general psd-rank problem which differs in our requirements for normalization: ρ_x should not only be psd but also trace 1, and E_{yz} should not only be psd but should satisfy $\sum_z E_{yz} = I$ for all y.

In this paper we study the question of when a quantum model can be compressed to a smaller dimension. We give conditions under which compression is, or is not, possible.

Main results. We present one theorem demonstrating *incompressibility* (Theorem 1) and three theorems describing *compressibility* (Theorems 2, 3 and 4). A common theme will be that measurement operators with high trace are a barrier to compressibility.

Theorem 1. *Let $\mathcal{D} \in \mathbb{R}^{X\times YZ}$, let $\left(E_{yz}\right)_{z=1}^{Z}$ be a d-dimensional POVMs and let ρ_x, $x \in [X]$, be d-dimensional quantum states. Assume that $\mathcal{D}_{x,yz} = \mathrm{tr}\left(\rho_x E_{yz}\right)$ for all $x \in [X]$, $y \in [Y]$ and $z \in [Z]$. Fix y. For all $z \in [Z]$, set $c_z^* = \max\{\mathcal{D}_{x;yz}\}_{x=1}^{X}$. Then, $d \geq \sum_{z=1}^{Z} c_z^*$. Moreover, if $\mathcal{D}' \in \mathbb{R}^{X\times YZ}$ is a set of observations coming from d'-dimensional states that satisfies $\|\mathcal{D} - \mathcal{D}'\|_\infty \leq \varepsilon$, then we have that $d' \geq -Z\varepsilon + \sum_{z=1}^{Z} c_z^*$.*

Theorem 1 is proven in section II. Theorem 1 was obtained independently in [14] (Theorem 24).

Note that the lower bound from Theorem 1 cannot exceed Z, i.e., the number of measurement outcomes. Consequently, Theorem 1 leaves open the possibility to compress quantum models of dimension $D > Z$ into models of dimension Z. This is indeed possible if for each measurement E_y all but one POVM element have trace norm constant in D. To approach this conclusion we first consider the compression of psd factorizations of $\mathcal{D}$.

Theorem 2. *Let $M_1, \dots, M_J$ be psd matrices on $\mathbb{C}^D$ let $\varepsilon \in (0, 1/2]$ and fix $d \in \mathbb{N}$ such that*

$$d > \frac{16}{\varepsilon^2} \ln(2JD). \tag{1}$$

Then there exist psd matrices M'_j on $\mathbb{C}^d$ such that for all $i, j \in [J]$

$$\begin{aligned} &\operatorname{tr}(M_i M_j) - \operatorname{tr}(M_i)\operatorname{tr}(M_j) 192\varepsilon \\ &\quad \leq \operatorname{tr}\big(M'_i M'_j\big) \leq \operatorname{tr}(M_i M_j) + \operatorname{tr}(M_i)\operatorname{tr}(M_j) 192\varepsilon. \end{aligned} \tag{2}$$

The matrices M'_j can be computed in randomized polynomial time in J and D.

Theorem 2 is proven in section III. It can be used to compress psd factorizations of $\mathcal{D}$ into approximate psd factorizations of $\mathcal{D}$. In particular, it can be used to compress quantum models for $\mathcal{D}$ into approximate psd factorizations of $\mathcal{D}$. However, we have no guarantee that the compression satisfies the normalization conditions which played the key role in the derivation of Theorem 1. This problem is addressed in Theorem 3.

Theorem 3. *For each $x \in [X]$, let ρ_x be a D-dimensional quantum state and for all $y \in [Y]$ let $\big(E_{yz}\big)_{z=1}^{Z}$ be a D-dimensional POVM. Set $J = X + YZ$. Let $\varepsilon \in (0, 1/2]$ and fix $d \in \mathbb{N}$ such that for all y*

$$\begin{aligned} d &> \frac{32}{\varepsilon^2} \ln(4JD) \\ d &> \frac{32}{\varepsilon^2} \operatorname{rank}\Big(\sum_{z=1}^{Z-1} E_{yz}\Big) \end{aligned} \tag{3}$$

Then, there exist d-dimensional quantum states ρ'_x and d-dimensional POVMs $\big(E'_{yz}\big)_{z=1}^{Z}$ such that

- *for all $x \in [X], y \in [Y], z \in [Z-1]$,*

$$\big|\operatorname{tr}\big(\rho'_x E'_{yz}\big) - \operatorname{tr}\big(\rho_x E_{yz}\big)\big| \leq 200\varepsilon \operatorname{tr}(E_{yz}).$$

- *For all $x \in [X], y \in [Y]$ and $z = Z$,*

$$\big|\operatorname{tr}\big(\rho'_x E'_{yz}\big) - \operatorname{tr}\big(\rho_x E_{yz}\big)\big| \leq 200\varepsilon \operatorname{tr}(I - E_{yZ}).$$

The compressed quantum model $\big(\rho'_x\big)_x$, $\big(E'_{yz}\big)_{yz}$ can be computed in randomized polynomial time in X, Y, Z and D.

Similar approximation guarantees hold for overlaps between pairs of states and pairs of measurements (see section IV C).

In the above theorem the outcome $z = Z$ is special, and our bounds are adapted to the case when E_{yZ} has much larger rank than the other POVM elements. The practical relevance of measurements of that kind is discussed in section IV E in the context of different applications of Theorem 3.

Theorem 3 is proven in section IV A to section IV D. Note that the main difference between the conditions in Theorem 3 and the conditions in Theorem 2 is the new rank-based constraint in (3). It is the direct consequence of the normalization conditions for quantum models. As an example consider one particular measurement E_y of the form $E_{yz} = |z\rangle\langle z|$ for $z \in [Z-1]$ and $E_{yZ} = I - \sum_{z=1}^{Z-1} E_{yz}$. Assume that the experimental states ρ_x are approximately equal to $|x\rangle\langle x|$. Then, by Theorem 1, $d \geq Z$. On the other hand, by Theorem 3, there exists an approximate quantum model with

$$d = \frac{32}{\varepsilon^2} \ln(JD) + \frac{32}{\varepsilon^2} Z \tag{4}$$

i.e., the lower bound is almost achieved.

Condition (3) is phrased in terms of the rank of POVM elements. This is the main caveat of Theorem 3 because in experimental systems we expect POVM elements to be full rank with light spectral tails. The following Theorem 4 is another version of Theorem 3 which tolerates POVM elements with high rank if their spectrum is decaying exponentially. See section V for its proof.

Theorem 4. *Let ρ_x, E_{yz} and J be as in Theorem 3. Denote by $\big(\varepsilon_n^{(y)}\big)_n$ the spectrum (ordered descendingly) of $\sum_{z=1}^{Z-1} E_{yz}$. To describe potentially thin spectral tails we assume that for some $j^*, b > 0$ and for all $j > 0$,*

$$\varepsilon_{j^*+j} \leq e^{-bj}. \tag{5}$$

Let $\varepsilon \in (0, 1/2]$ and fix $d \in \mathbb{N}$ such that

$$d \geq \frac{128}{\varepsilon^2} \ln(4JD) \tag{6}$$

$$d \geq \frac{128}{\varepsilon^2} \left(j^* + \frac{1}{b} \ln \frac{8}{\varepsilon}\right) \tag{7}$$

Then, there exist d-dimensional quantum states ρ'_x and d-dimensional POVMs $\big(E'_{yz}\big)_{z=1}^{Z}$ with the same approximation promises as Theorem 3. Again, the compressed quantum model $\big(\rho'_x\big)_x$, $\big(E'_{yz}\big)_{yz}$ can be computed in randomized polynomial time in X, Y, Z and D.

Implications. Bounding the approximate psd-rank of identity matrices is a problem addressed in [14] (section 6.2) and a lower bound was derived. Theorem 2 on the other hand provides upper bounds for general matrices.

Theorem 3 and 4 have the potential to alter our perspective in quantum tomography in that it reveals an equivalence between high-dimensional models and low-dimensional models if the measurements are low rank and have few outcomes. This is discussed in section IV E 1.

Another lesson to be learned here is that whether we interpret measurement outcomes as part of a single measurement or whether we summarize small groups of measurement outcomes to individual measurements hugely affects the compressibility of the model. To illustrate this point, let $\{|j\rangle\}_{j=1}^{D}$ denote an orthonormal basis in the Hilbert space $\mathbb{C}^D$, let $X = D$ with $\rho_x = |x\rangle\langle x|$, and let $Y = 1$, $Z = D$ with $E_{1z} = |z\rangle\langle z|$. It follows that $\big(\mathrm{tr}(\rho_x E_{yz})\big)_{xyz}$ equals the identity matrix of size D. Then, by Theorem 1, the dimension of each model describing $\big(\mathrm{tr}(\rho_x E_{yz})\big)_{xyz}$ is lower bounded by D and cannot be compressed. Even allowing an error of $\leq \delta$ in each entry cannot change the lower bound by more than $D\delta$ which implies that the dimension must be $\geq (1-\delta)D$. Indeed this demonstrates that the dependence of d on the rank of the sum of the POVM operators in Theorem 3 cannot in general be removed. On the other hand, if instead we performed D individual measurements $\big(|z\rangle\langle z|, I - |z\rangle\langle z|\big)$ to test for each of the D possible inputs then the model can be compressed exponentially. Consequently, it becomes impossible to resolve any high complexity of the considered physical system.

The two models described in the last paragraph correspond to the communication tasks of sending and identifying, respectively, a $\log(D)$-bit classical message. For these models Winter's 2004 paper [17] observed the same exponential compression that we describe. Our work can be thought of as a generalization of [17] in that it replaces the simulation of orthonormal states and binary measurements $\{|z\rangle\langle z|, I - |z\rangle\langle z|\}$ with the simulation of general states and measurements. It is qualitatively similar in that it works best when compressing low-rank measurements.

Further applications of Theorem 3 involve a Corollary of Theorem 3 (see Section IV E 3) that allows for shrinkage of upper bounds on the one-way quantum communication complexity $Q^1(f)$ if the upper bound is based on high-rank measurements. Moreover, we discuss fundamental limitations on robust dimension witnessing [1, 4, 7, 10, 15, 16] in Section IV E 2.

II. INCOMPRESSIBILITY

Consider a D-dimensional quantum model for $\mathcal{D}$. Due to normalization of measurements, $\sum_z E_{yz} = I$. On a D-dimensional Hilbert space the identity matrix has trace-norm D. Therefore, if $\mathcal{D}$ implies lower bounds on $\|I\|_1$ then these bounds can be used as lower bounds on D. Indeed the derivation of lower bounds on $\|I\|_1$ using $\mathcal{D}$ is straightforward: Consider data $\mathcal{D} \in \mathbb{R}^{X\times Z}$ generated by quantum states $(\rho_x)_{x=1}^{X}$ and a single POVM $(E_z)_{z=1}^{Z}$ on $\mathbb{C}^d$, i.e., $\mathcal{D}_{xz} = \mathrm{tr}\big(\rho_x E_z\big)$. For each $z \in [Z]$ let $c_z^* := \max\{\mathcal{D}_{xz}\}_{x=1}^{X}$ and $x' := \mathrm{argmax}\{\mathcal{D}_{xz}\}_{x=1}^{X}$. Then,

$$c_z^* = \mathrm{tr}\big(\rho_{x'} E_z\big) \leq \|\rho_{x'}\|\|E_z\|_1$$

and therefore, $\|E_z\|_1 \geq c_z^*$ because $\|\rho\| \leq 1$ for all quantum states ρ. Here, $\|\cdot\|$ denotes the operator norm (i.e. largest singular value) and $\|\cdot\|_1$ the trace norm (i.e. sum of singular values). It follows that for all $z \in [Z]$, $\mathrm{tr}(E_z) \geq c_z^*$ and consequently,

$$\mathrm{tr}\Big(\sum_{z=1}^{Z} E_z\Big) \geq \sum_{z=1}^{Z} c_z^*.$$

By assumption $\big(E_z\big)_{z=1}^{Z}$ is a POVM. So in particular, $\sum_{z=1}^{Z} E_z = I_d$ and therefore,

$$d = \mathrm{tr}(I_d) = \mathrm{tr}\Big(\sum_{z=1}^{Z} E_z\Big) \geq \sum_{z=1}^{Z} c_z^*. \tag{8}$$

This is a lower bound on the dimension of any quantum model describing $\mathcal{D}$. It has been found independently in [14] (Theorem 24).

It is straightforward to see how the dimension lower bound $l(\mathcal{D}) := \sum_{z=1}^{Z} c_z^*$ reacts to noise in $\mathcal{D}$. Assume that $\mathcal{D}'$ is a noisy approximation of $\mathcal{D}$ in the sense that $\|\mathcal{D} - \mathcal{D}'\|_\infty \leq \varepsilon$. Here, $\|\mathcal{D} - \mathcal{D}'\|$ denotes the maximum norm associated to the vectorization of $\mathcal{D} - \mathcal{D}'$. Then,

$$\begin{aligned}
&|l(\mathcal{D}) - l(\mathcal{D}')| \\
&\leq \sum_{z=1}^{Z} \Big| \max\{\mathcal{D}_{xz}\}_{x=1}^{X} - \max\{\mathcal{D}'_{xz}\}_{x=1}^{X} \Big| \\
&= \sum_{z=1}^{Z} \max\Big\{ \max\{\mathcal{D}_{xz}\}_{x=1}^{X} - \max\{\mathcal{D}'_{xz}\}_{x=1}^{X}, \\
&\qquad\qquad \max\{\mathcal{D}'_{xz}\}_{x=1}^{X} - \max\{\mathcal{D}_{xz}\}_{x=1}^{X} \Big\} \\
&\leq \sum_{z=1}^{Z} \max\Big\{ \max\{\mathcal{D}'_{xz} + \varepsilon\}_{x=1}^{X} - \max\{\mathcal{D}'_{xz}\}_{x=1}^{X}, \\
&\qquad\qquad \max\{\mathcal{D}_{xz} + \varepsilon\}_{x=1}^{X} - \max\{\mathcal{D}_{xz}\}_{x=1}^{X} \Big\} \\
&= Z\varepsilon.
\end{aligned}$$

In other words, if we measure $\mathcal{D}'$ with accuracy $\|\mathcal{D} - \mathcal{D}'\|_\infty \leq \varepsilon$, then we know that $d \geq l(\mathcal{D}') - Z\varepsilon$.

As an example, we consider data generated by $E_z = |z\rangle\langle z|$, $\rho_z = |\pi(z)\rangle\langle\pi(z)|$ with $z \in [D]$, $\pi \in S_D$ a permutation and with $(|z\rangle)_{z=1}^{D}$ being an orthonormal basis in $\mathbb{C}^D$. It follows that for all $z \in [D]$, $c_z^* = 1$. By (8), $d \geq D$, i.e., the considered quantum system generating the considered data cannot be compressed into a lower-dimensional quantum system. This example generalizes in obvious manners: if for each measurement outcome there exists a state that can be measured with approximate certainty, then roughly, the considered quantum system cannot be compressed into another quantum system whose dimension exceeds Z.

We conclude that general D-dimensional quantum models cannot be compressed into d-dimensional quantum models with $d \ll D$. However, this still leaves room for the existence of compression schemes that can compress specific classes of models. This is what we are going to explore next.

III. COMPRESSION OF POSITIVE SEMIDEFINITE FACTORIZATION

Let $\mathcal{M} \in \mathbb{R}_+^{N\times M}$, let $J := N + M$ and let $S^+(\mathbb{C}^D)$ denote the set of psd matrices on $\mathbb{C}^D$. The psd matrices $\big(A_n\big)_{n=1}^N$ and $\big(B_m\big)_{m=1}^M$ from $S^+(\mathbb{C}^D)$ are said to provide a D-dimensional psd factorization of $\mathcal{M}$ if for all entries $\mathcal{M}_{nm} = \mathrm{tr}\big(A_n B_m\big)$. The *psd-rank* of $\mathcal{M}$ is defined to be the dimension of the lowest-dimensional psd factorization of $\mathcal{M}$. See [5] for a recent review of psd-rank.

A. The compression scheme

Assume $\big(A_n\big)_{n=1}^N$ and $\big(B_m\big)_{m=1}^M$ form a D-dimensional psd factorization of $\mathcal{M}$ and set $M_n = A_n$ for $1 \le n \le N$ and $M_{N+n} = B_n$ for $1 \le n \le M$. Now and in the remainder we assume that $\Pi = G/\sqrt{2d} \in \mathbb{C}^{d\times D}$ with $G_{ij} = S_{ij} + iT_{ij}$ where $S_{ij}, T_{ij} \sim \mathcal{N}(0,1)$ independent and identically distributed (iid). This normalization is chosen so that $\mathbb{E}[\Pi] = 0$ and $\mathbb{E}[\Pi^\dagger\Pi] = I_D$. In this section we show that for all $n, m \in [J]$, the map

$$M_j \mapsto \Pi M_j \Pi^\dagger \tag{9}$$

has the property that

$$\mathrm{tr}\big(\Pi M_n \Pi^\dagger \Pi M_m \Pi^\dagger\big) \approx \mathrm{tr}\big(M_n M_m\big). \tag{10}$$

Thus, it approximately preserves all of the Gram matrix (i.e., not just $\mathcal{M}$) associated to the considered psd factorization. Since $d < D$, the map (9) can be regarded as a *compression* of the psd factorization we started with. The compression errors are additive. The additive error associated to $\mathrm{tr}(A_n B_m)$ scales with $\mathrm{tr}\big(A_n\big)\mathrm{tr}\big(B_m\big)$.

This proof is divided into two pieces. First we argue that (10) holds for rank-1 matrices M_m, M_n in section III B and then we extend to general psd matrices M_m, M_n in section III C.

B. Compression errors for vectors

The errors of the compression of pure states correspond to the famous Johnson-Lindenstrauss Lemma, which has the following formulation for complex vector spaces:

Theorem 5 (complex Johnson-Lindenstrauss [12])**.** *Assume that $\Pi = G/\sqrt{2d} \in \mathbb{C}^{d\times D}$ with $G_{ij} = A_{ij} + iB_{ij}$ where $A_{ij}, B_{ij} \sim \mathcal{N}(0,1)$ iid. Let $v_1, ..., v_S$ be arbitrary vectors in $\mathbb{C}^D$ and $\varepsilon \in (0,1)$. Then,*

$$\mathbb{P}\big[\forall i\ \|\Pi v_i\|_2 \in [1-\varepsilon, 1+\varepsilon]\|v_i\|_2\big] \ge 1 - 2Se^{d\varepsilon^2/8} \tag{11}$$

The proof is simple enough that we reproduce it here.

Proof. Fix a particular $v := v_i$. By linearity of Π, we can assume without loss of generality that $\|v\|_2 = 1$. Let U denote a unitary matrix with the property $v = Ue_1$, where e_1 denotes the vector with a one in the first position and zeroes elsewhere. By unitary invariance of the complex Gaussian measure on $\mathbb{C}^D$, $\Pi \sim \Pi U$ (i.e. the two random variables are identically distributed) and so

$$\|\Pi v\|_2^2 = \|\Pi U e_1\|_2^2 \sim \|\Pi e_1\|_2^2 = \sum_{j=1}^{d} \frac{A_{1j}^2 + B_{1j}^2}{2d} \tag{12}$$

This average of the square of $2d$ standard normal random variables is known as the χ^2 distribution, and its concentration of measure properties are standard. Indeed, by Corollary 5.5 in [2],

$$\mathbb{P}\big[\|\Pi v\|_2^2 \ge 1 + 2\varepsilon\big] \le e^{-\frac{1}{2}\varepsilon^2 d}$$
$$\mathbb{P}\big[\|\Pi v\|_2^2 \le 1 - \varepsilon\big] \le e^{-\frac{1}{2}\varepsilon^2 d}$$

so that

$$\mathbb{P}\big[\|\Pi z\|_2^2 \in [1-\varepsilon, 1+\varepsilon]\big] \ge 1 - 2e^{-\frac{1}{8}\varepsilon^2 d}.$$

Eq. (11) follows by the union bound:

$$\mathbb{P}\Big[\bigcup_{i\in[S]} \Big\{\|\Pi v_i\|_2^2 \notin [(1-\varepsilon)\|v_i\|^2, (1+\varepsilon)\|v_i\|^2]\Big\}\Big]$$
$$\le S \cdot 2e^{-\frac{1}{8}\varepsilon^2 d}$$

□

The usual Johnson-Lindenstrauss Lemma shows that Π preserves (with high probability) not only the lengths of a collection of vectors, but also their pairwise distances. In fact, we will demand slightly more.

Corollary 6. *Let $v_1, \ldots, v_S$, Π, ε be as in Theorem 5. Then with probabilty $\ge 1 - 4S^2 e^{-d\varepsilon^2/8}$ we have*

$$\|\Pi v_i\|_2 \in [1-\varepsilon, 1+\varepsilon]\|v_i\|_2 \tag{13a}$$
$$\|\Pi(v_i + v_j)\|_2 \in [1-\varepsilon, 1+\varepsilon]\|v_i + v_j\|_2 \tag{13b}$$
$$\|\Pi(v_i + iv_j)\|_2 \in [1-\varepsilon, 1+\varepsilon]\|v_i + iv_j\|_2 \tag{13c}$$
$$\|\Pi(v_i - v_j)\|_2 \in [1-\varepsilon, 1+\varepsilon]\|v_i - v_j\|_2 \tag{13d}$$
$$\|\Pi(v_i - iv_j)\|_2 \in [1-\varepsilon, 1+\varepsilon]\|v_i - iv_j\|_2 \tag{13e}$$

Proof. Apply Theorem 5 to the $S + 4\binom{S}{2} \le 2S^2$ vectors $\{v_i\}_{i\in[S]} \cup \{v_i \pm v_j, v_i \pm iv_j\}_{1\le i<j\le S}$. □

Matrices Π satisfying (13) for some set of vectors $\{v_1, \ldots, v_S\}$ are said to be ε-JL matrices.

C. Compression errors for positive semidefinite matrices

In this section we extend Theorem 5 to show that general psd matrices have their inner products approximately preserved by a random compression map.

First we show that ε-JL matrices also approximately preserve inner products between rank-1 psd matrices (recall that $\mathrm{tr}(|v\rangle\langle v||w\rangle\langle w|) = |\langle v|w\rangle|^2$).

Lemma 7. *If Π is an ε-JL matrix for the set $\{v_1, \ldots, v_S\}$ (i.e. satisfies (13)), then*

$$\big|\,|\langle v_i|\Pi^\dagger\Pi|v_j\rangle|^2 - |\langle v_i|v_j\rangle|^2\big| \leq 192\varepsilon\|v_i\|_2^2\|v_j\|_2^2. \tag{14}$$

While a direct analysis of the Gaussian probability distribution would yield a sharper constant, our approach highlights the fact that it is only the ε-JL property that is needed.

Proof. Since (14) is homogenous, we can assume that v_i, v_j are unit vectors. By the polarization identity,

$$\langle v_i|\Pi^\dagger\Pi|v_j\rangle = \frac{1}{4}\Big(x_1 - x_2 + ix_3 - ix_4)\Big) \tag{15}$$

with

$$\begin{aligned} x_1 &= \big\|\Pi(v_i + v_j)\big\|_2^2 & x_2 &= \big\|\Pi(v_i - v_j)\big\|_2^2 \\ x_3 &= \big\|\Pi(v_i + iv_j)\big\|_2^2 & x_4 &= \big\|\Pi(v_i - iv_j)\big\|_2^2 \end{aligned}$$

Thus

$$|\langle v_i|\Pi^\dagger\Pi|v_j\rangle|^2 = \frac{1}{16}\Big((x_1 - x_2)^2 + (x_3 - x_4)^2\Big) =: f(\vec{x}) \tag{16}$$

Define as well

$$\begin{aligned} y_1 &= \big\|v_i + v_j\big\|_2^2 & y_2 &= \big\|v_i - v_j\big\|_2^2 \\ y_3 &= \big\|v_i + iv_j\big\|_2^2 & y_4 &= \big\|v_i - iv_j\big\|_2^2 \end{aligned}$$

Observe that each $y_i \leq 4$ and that

$$\langle v_i|v_j\rangle = \frac{1}{4}\Big(y_1 - y_2 + iy_3 - iy_4)\Big) \tag{17}$$

By the ε-JL property of Π, we have that for $j \in [4]$, $x_j \in [(1-\varepsilon)^2, (1+\varepsilon)^2]y_j$. Set

$$\begin{aligned} K :=& [(1-\varepsilon)^2 y_1, (1+\varepsilon)^2 y_1] \times [(1-\varepsilon)^2 y_2, (1+\varepsilon)^2 y_2] \times \\ & [(1-\varepsilon)^2 y_3, (1+\varepsilon)^2 y_3] \times [(1-\varepsilon)^2 y_4, (1+\varepsilon)^2 y_4]. \end{aligned}$$

We are going to apply

$$|f(\vec{x}) - f(\vec{y})| \leq \|\nabla f\|_K \|\vec{x} - \vec{y}\|_\infty \tag{18}$$

where $\|\nabla f\|_K = \max\big\{\|\nabla f\|_1 \,|\, \vec{x} \in K\big\}$. By definition of $f(\vec{x})$,

$$\begin{aligned} \|\nabla f\|_K &= \frac{1}{4}\max_{\vec{x}\in K}\big\{|x_1 - x_2| + |x_3 - x_4|\big\} \\ &\leq \frac{1}{4}\max_{\vec{x}\in K}\big\{|x_1| + |x_2| + |x_3| + |x_4|\big\} \\ &\leq \frac{(1+\varepsilon)^2}{4}\big\{|y_1| + |y_2| + |y_3| + |y_4|\big\} \\ &\leq 4(1+\varepsilon)^2 \leq 16, \end{aligned} \tag{19}$$

using $\varepsilon \leq 1$ in the last step. We also bound

$$\begin{aligned} \|\vec{x} - \vec{y}\|_\infty &= \max_j |x_j - y_j| \\ &\leq (2\varepsilon + \varepsilon^2)\max_j y_j \\ &\leq 12\varepsilon. \end{aligned} \tag{20}$$

Thus by Eq. (18),

$$|f(\vec{x}) - f(\vec{y})| \leq 192\varepsilon \tag{21}$$

□

Armed with this fact, we can prove (10).

Theorem 2 (restatement). *Let $M_1, \ldots, M_J$ be psd matrices on $\mathbb{C}^D$ let $\varepsilon \in (0, 1/2]$ and fix $d \in \mathbb{N}$ such that*

$$d > \frac{16}{\varepsilon^2}\ln\big(2JD\big). \tag{22}$$

Then there exist psd matrices M_j' on $\mathbb{C}^d$ such that for all $i, j \in [J]$

$$\begin{aligned} &\mathrm{tr}(M_i M_j) - \mathrm{tr}(M_i)\mathrm{tr}(M_j)192\varepsilon \\ &\quad\leq \mathrm{tr}\big(M_i' M_j'\big) \leq \mathrm{tr}(M_i M_j) + \mathrm{tr}(M_i)\mathrm{tr}(M_j)192\varepsilon. \end{aligned} \tag{23}$$

The matrices M_j' can be computed in randomized polynomial time in J and D.

Proof. For each $i \in [J]$, let the eigendecomposition of M_i be

$$M_i = \sum_{a=1}^{D} \lambda_a^i |\psi_a^i\rangle\langle\psi_a^i|.$$

Observe that $\|M_i\|_1 = \sum_{a=1}^D \lambda_a^i$.

Choose $\Pi \in \mathbb{C}^{d\times D}$ according to a Gaussian distribution as in Theorem 5 (i.e. such that $\mathbb{E}[\Pi] = 0$ and $\mathbb{E}[\Pi^\dagger\Pi] = I_D$). By Corollary 6, Π is an ε-JL matrix for the vectors $\{|\psi_a^i\rangle\}$ with probability

$$\geq 1 - 4J^2D^2e^{-d\varepsilon^2/8}. \tag{24}$$

By (22), this is > 0, thus implying that there exists some Π with this property. Fix this choice of Π for the rest of the proof.

Define

$$M_i' := \Pi M_i \Pi^\dagger. \tag{25}$$

We now compute

$$\begin{aligned}\mathrm{tr}(M_i' M_j') &= \mathrm{tr}(\Pi M_i \Pi^\dagger \Pi M_j \Pi^\dagger) \\ &= \sum_{a,b=1}^{D} \lambda_a^i \lambda_b^j |\langle \psi_a^i | \Pi^\dagger \Pi | \psi_b^j \rangle|^2\end{aligned}$$

By Lemma 7, we have

$$\left| |\langle \psi_a^i | \Pi^\dagger \Pi | \psi_b^j \rangle|^2 - |\langle \psi_a^i | \psi_b^j \rangle|^2 \right| \leq 192\varepsilon, \tag{26}$$

for each i, j, a, b. Thus

$$\begin{aligned}|\mathrm{tr}(M_i' M_j') - \mathrm{tr}(M_i M_j)| &\leq \sum_{a,b=1}^{D} \lambda_a^i \lambda_b^j 192\varepsilon \\ &= \|M_i\|_1 \|M_j\|_1 192\varepsilon \end{aligned} \tag{27}$$

□

We remark that this proof would work equally well with general matrices, with the error terms scaling as $\|M_i\|_1 \|M_j\|_1$ instead of $\mathrm{tr}(M_i)\mathrm{tr}(M_j)$. However, if we are willing to abandon the psd condition, a much tighter bound is possible simply by treating $\{M_i\}$ as vectors in $\mathbb{C}^{d^2}$ and using Theorem 5. This would result in errors proportional instead to $\|M_i - M_j\|_2^2$ (or optionally $\|M_i\|_2 \|M_j\|_2$), both of which can be smaller than $\|M_i\|_1 \|M_j\|_1$ by as much as D.

More generally, one might ask whether a nonlinear compression map can remove this dependence on the matrix norm. However, the absolute error cannot be invariant under $M_n \mapsto \lambda M_n$ ($\lambda > 0$). This can be understood as follows. Assume that $\hat{\mathcal{C}}$ is a compression with the property that there exist functions $c(\varepsilon)$ and $d(\varepsilon)$ such that

$$\mathrm{tr}(\hat{\mathcal{C}}(M_n)\hat{\mathcal{C}}(M_m)) = c(\varepsilon)\mathrm{tr}(M_n M_m) + d(\varepsilon)$$

for all positive semidefinite matrices M_n, M_m and which are invariant under $M_n \mapsto \lambda M_n$. Then,

$$\mathrm{tr}(\hat{\mathcal{C}}(\lambda |n\rangle\langle n|)\hat{\mathcal{C}}(\lambda |m\rangle\langle m|)) = c(\varepsilon)\lambda^2 \delta_{nm} + d(\varepsilon)$$

where $(|n\rangle)_n$ denotes the canonical basis in $\mathbb{C}^D$. Set $\eta_n(\lambda) := \sqrt{c(\varepsilon)\lambda^2 + d(\varepsilon)}$. It follows that

$$\left\| \frac{1}{\eta_n} \hat{\mathcal{C}}(\lambda |n\rangle\langle n|) \right\|_2 = 1$$

and for $m \neq n$

$$\mathrm{tr}\left(\frac{1}{\eta_n} \hat{\mathcal{C}}(\lambda |n\rangle\langle n|) \frac{1}{\eta_m} \hat{\mathcal{C}}(\lambda |m\rangle\langle m|) \right) = \frac{d(\varepsilon)}{\eta_n \eta_m} \to 0$$

as $\lambda \to \infty$ because, by assumption, $d(\varepsilon)$ is invariant under $M_n \mapsto \lambda M_n$. Assuming $D > d^2$, we would be able to construct an overcomplete orthonormal basis of positive semidefinite matrices. This is impossible and we conclude that necessarily, $d(\varepsilon) = \Omega(\lambda^2)$.

IV. COMPRESSION OF QUANTUM MODELS WITH FEW MEASUREMENT OUTCOMES

The purpose of this section is the definition and the analysis of a scheme to compress quantum models. The discussion in section II shows that the compressibility of quantum models is limited by the number of measurement outcomes. This will be reflected in the conditions for the compression to be applicable.

Hence, given $\mathcal{D}$ generated by arbitrary quantum states ρ_x and POVMs $\left(E_{yz}\right)_{z=1}^{Z}$ on $\mathbb{C}^D$, our goal is to find low-dimensional states and POVMs on $\mathbb{C}^d$ which reproduce $\mathcal{D}$ approximately ($d < D$). To treat the states and the POVMs all at once, we define $M_x := \rho_x$ for $x \in [X]$ and $M_{X+(y-1)Z+z} := E_{yz}$ for $y \in [Y]$ and $z \in [Z]$. So the matrices $(M_n)_{n=1}^{J}$ form a list of matrices describing the states and POVMs. Then, the matrix $\mathcal{G}_{nm} = \mathrm{tr}\left(M_n M_m\right)$ is the Gram matrix generated by all of the states and all of the measurement operators. Hence, $\mathcal{G}$ describes the Euclidean geometry of the set $\left(M_n\right)_{n=1}^{J}$.

The transition from the high-dimensional states and POVMs to their low-dimensional counterparts will be described in terms of a compression map $\mathcal{C} : S^+(\mathbb{C}^D)^{\times J} \to S^+(\mathbb{C}^d)^{\times J}$,

$$\mathcal{C} : \left(M_n\right)_{n=1}^{J} \mapsto \mathcal{C}\Big(\left(M_n\right)_{n=1}^{J}\Big) = \left(\mathcal{C}_n(M_n)\right)_{n=1}^{J}.$$

We will see that the proposed compression scheme $\mathcal{C}$ approximately preserves not only the inner products $\mathcal{M}_{n,m}$ but also the entire Gram matrix $\mathcal{G}$. The compression errors are additive. They scale with the trace norm of the measurement operators. In the remainder we are going to suppress the index n of $\mathcal{C}_n(\cdot)$.

A. The compression scheme

Let $\varepsilon > 0$ and set $J = X + YZ$. Assume that $\Pi = G/\sqrt{2d} \in \mathbb{C}^{d \times D}$ with $G_{ij} = A_{ij} + iB_{ij}$ where $A_{ij}, B_{ij} \sim \mathcal{N}(0,1)$ iid. Set

$$\mathcal{C}(\rho_x) = \frac{1}{\mathrm{tr}(\Pi \rho_x \Pi^\dagger)} \Pi \rho_x \Pi^\dagger \tag{28}$$

for all $x \in [X]$,

$$\mathcal{C}(E_{yz}) = \frac{1}{1+\varepsilon} \Pi E_{yz} \Pi^\dagger \tag{29}$$

for all $y \in [Y]$ and $z \in [Z-1]$ and set

$$\mathcal{C}(E_{yZ}) = I - \sum_{z=1}^{Z-1} \mathcal{C}(E_{yz}) \tag{30}$$

for all $y \in [Y]$. Note that the map $\mathcal{C}$ is non-linear. Obviously, $\mathcal{C}(\rho_x)$ are valid quantum states. In section IV B, we are going to show that with non-vanishing probability, the matrices $\left(\mathcal{C}(E_{yz})\right)_{z=1}^{Z}$ form a valid POVM and therefore, the image of $\mathcal{C}$ gives valid quantum states and measurements on $\mathbb{C}^d$.

B. Probability that compression of a POVM is psd and normalized

It is not *a priori* obvious that $\big(\mathcal{C}(E_{yz})\big)_{z=1}^{Z}$ is indeed a valid POVM. While $\mathcal{C}(E_{yz}) \geq 0$ for $z < Z$ and $\sum_{z=1}^{Z}\mathcal{C}(E_{yz}) = I$ hold automatically, it is not always true that

$$\mathcal{C}(E_{yZ}) \geq 0 \quad \Leftrightarrow \quad \Pi \sum_{z=1}^{Z-1} E_{yz}\Pi^\dagger \leq (1+\varepsilon)I \tag{31}$$

However, for appropriate choice of ε, we will see that (31) holds with high probability. Let P be a projector onto the support of $\sum_{z=1}^{Z-1} E_{yz}$. Let $r = \mathrm{rank}P = \mathrm{rank}\sum_{z=1}^{Z-1} E_{yz}$. Then $\|\Pi P \Pi^\dagger\| = \sigma_{\max}(\Pi P)^2$, where $\sigma_{\max}(\cdot)$ denotes the largest singular value. Because of the unitary invariance of the Gaussian measure, $A := \Pi P$ is distributed identically to a $d \times r$ matrix of complex i.i.d. Gaussians with mean 0 and variance $1/d$. Specifically observe that $\mathbb{E}[AA^\dagger] = I_r$.

We now appeal to a standard result in random matrix theory.

Theorem 8 ([9]). *Let A be a $r \times d$ Gaussian matrix with iid entries satisfying $\mathbb{E}[A_{ij}] = 0, \mathbb{E}[|A_{ij}|^2] = 1/d$. Then for $0 \leq t \leq d/2$ and $0 \leq \delta \leq 2$ we have*

$$\mathbb{E}[\exp(t\|AA^\dagger\|)] \leq d\exp\left(t(1+\sqrt{r/d})^2 + \frac{t^2}{d}(1+r/d)\right) \tag{32}$$

$$\mathbb{P}\big[\|AA^\dagger\| \geq (1+\sqrt{r/d})^2 + \delta\big] \leq de^{-d\delta^2/8} \tag{33}$$

Proof. (32) is Lemma 7.2 of [9] and (33) follows by setting $t = d\delta/4$, and using the bound $\mathbb{P}\big[\|AA^\dagger\| \geq \lambda\big] \leq \mathbb{E}[\exp(t\|AA^\dagger\|)]e^{-t\lambda}$ and the fact that $r \leq d$. □

This implies that our dimension-reduction strategy does not significantly blow up the POVM normalization as long as the rank of the measurement operators is much less than d.

Lemma 9. *Let $\Pi = G/\sqrt{2d} \in \mathbb{C}^{d\times D}$ with $G_{ij} = A_{ij} + iB_{ij}$ where $A_{ij}, B_{ij} \sim \mathcal{N}(0,1)$ iid. Assume $E \in \mathbb{C}^{D\times D}$ psd with $\|E\| \leq 1$. Then,*

$$\mathbb{P}\big[\|\Pi E \Pi^\dagger\| \leq 1+\varepsilon\big] \geq 1 - de^{-\frac{\varepsilon^2}{32}d} \tag{34}$$

if

$$d \geq \frac{32}{\varepsilon^2}\,\mathrm{rank}(E). \tag{35}$$

Proof. Let $r = \mathrm{rank}E$. Since $E \leq I$, then $E \leq P$ for some rank-r projector P. Then, according to the discussion earlier in this section, $\|\Pi E\Pi^\dagger\| \leq \|AA^\dagger\|$ for some matrix A satisfying the conditions of Theorem 8. If (66) holds then with probability $\geq 1 - de^{-d\varepsilon^2/32}$ we have

$$\|\Pi E\Pi^\dagger\| \leq \|AA^\dagger\| \leq (1+\sqrt{r/d})^2 + \varepsilon/2. \tag{36}$$

From (66) and $\varepsilon \leq 1/2$ we have

$$\begin{aligned}(1+\sqrt{r/d})^2 &\leq (1+\varepsilon/\sqrt{32})^2 \\ &\leq 1 + \varepsilon\left(\frac{2}{\sqrt{32}} + \frac{1}{64}\right) \leq 1 + 0.37\varepsilon.\end{aligned} \tag{37}$$

Together this implies that $\|\Pi E\Pi^\dagger\| \leq 1+\varepsilon$. □

We now apply this to dimension reduction. For each $y \in [Y]$ set

$$\mathcal{E}_y = \{\Pi \mid \|\Pi\Big(\sum_{z\in[Z-1]} E_{yz}\Big)\Pi^\dagger\| \leq 1+\varepsilon\}. \tag{38}$$

By Lemma 9 and the union bound

$$\mathbb{P}\big[\bigcap_{y\in[Y]} \mathcal{E}_y\big] \geq 1 - Ye^{-\frac{\varepsilon^2}{32}d} \tag{39}$$

if

$$d > \frac{32}{\varepsilon^2}\,\mathrm{rank}\Big(\sum_{z=1}^{Z-1} E_{yz}\Big). \tag{40}$$

C. Tail bounds on compression errors

We use (23) to separately list the compression errors for inner products $\mathrm{tr}\big[\rho_x\rho_{x'}\big]$, $\mathrm{tr}\big[\rho_x E_{yz}\big]$ and $\mathrm{tr}\big[E_{yz}E_{y'z'}\big]$. Let $\mathcal{J}$ denote the event where Π satisfies the ε-JL property for the eigenstates of all the ρ_x and E_{yz}.

This event implies the following bounds.

$$\begin{aligned}\mathrm{tr}\big[\Pi\rho_x\Pi^\dagger\big] &= \sum_{j=1}^{D} p_j \mathrm{tr}\big[\Pi|\psi_j^x\rangle\langle\psi_j^x|\Pi^\dagger\big] = \sum_{j=1}^{D} p_j \|\Pi|\psi_j^x\rangle\|_2^2 \\ &\in [(1-\varepsilon)^2, (1+\varepsilon)^2].\end{aligned} \tag{41}$$

State-state error. For $n, m \in [X]$, (23) implies (using (41))

$$\begin{aligned}&\frac{1}{(1+\varepsilon)^4}\Big(\mathrm{tr}(\rho_x\rho_{x'}) - 192\varepsilon\Big) \\ &\leq \mathrm{tr}\big(\mathcal{C}(\rho_x)\mathcal{C}(\rho_{x'})\big) \\ &\leq \frac{1}{(1-\varepsilon)^4}\Big(\mathrm{tr}(\rho_x\rho_{x'}) + 192\varepsilon\Big)\end{aligned} \tag{42}$$

State-measurement error. Let $z \in [Z-1]$. By (23),

$$\begin{aligned}&\frac{1}{(1+\varepsilon)^3}\Big(\mathrm{tr}(\rho_x E_{yz}) - \mathrm{tr}(E_{yz})192\varepsilon\Big) \\ &\leq \mathrm{tr}\big(\mathcal{C}(\rho_x)\mathcal{C}(E_{yz})\big) \\ &\leq \frac{1}{(1-\varepsilon)^2(1+\varepsilon)}\Big(\mathrm{tr}(\rho_x E_{yz}) + \mathrm{tr}(E_{yz})192\varepsilon\Big)\end{aligned} \tag{43}$$

Let $z = Z$. By (28), (29) and (30),

$$\begin{aligned}&\operatorname{tr}\big[\mathcal{C}(\rho_x)\mathcal{C}(E_{yZ})\big]\\&= \operatorname{tr}\Big[\frac{\Pi\rho_x\Pi^\dagger}{\operatorname{tr}[\Pi\rho_x\Pi^\dagger]}\Big(I - \Pi\Big(\sum_{z=1}^{Z-1} E_{yz}\Big)\Pi^\dagger\Big)\Big]\\&= 1 - \Big(\sum_{z=1}^{Z-1}\operatorname{tr}\big[\mathcal{C}(\rho_x)\mathcal{C}(E_{yz})\big]\Big)\end{aligned} \tag{44}$$

All the summands in (44) can be bounded by (43).

Measurement-measurement error. Let $z, z' \in [Z-1]$. By (23),

$$\begin{aligned}&\frac{1}{(1+\varepsilon)^2}\Big(\operatorname{tr}(E_{yz}E_{y'z'}) - \operatorname{tr}(E_{yz})\operatorname{tr}(E_{y'z'})192\varepsilon\Big)\\&\leq \operatorname{tr}\big(\mathcal{C}(\operatorname{tr}(E_{vk}))\mathcal{C}(E_{y'z'})\big)\\&\leq \frac{1}{(1+\varepsilon)^2}\Big(\operatorname{tr}(E_{vk}E_{y'z'}) + \operatorname{tr}(E_{vk})\operatorname{tr}(E_{y'z'})192\varepsilon\Big)\end{aligned} \tag{45}$$

For $z = Z$ we get bounds similar to (44).

In each case the denominators can be absorbed by rounding up 192ε to 200ε. This is a straightforward calculation in the case when $\varepsilon \leq 1/200$ and when $\varepsilon > 1/200$ the error bound is vacuously true.

Here, we are able to avoid the no-go Theorem from [11] (already for the factorization of the identity matrix) because $\mathcal{C}$ is not completely positive, or even linear. Indeed, generically each row of Π will have norm $\sqrt{D/d}$ and so we will have $\|\Pi\| \approx \sqrt{D/d}$ with high probability. Thus there will almost always exist matrices E such that $\|\Pi E\Pi^\dagger\| \gg \|E\|$. However, our compression scheme can work because $\|\Pi E\Pi^\dagger\|$ is small for most E (or more precisely, for all E it is small for most Π). This feature of being able to achieve something for a small number of vectors that is impossible to extend to all vectors can be seen already in the original Johnson-Lindenstrauss Lemma.

D. Proof of Theorem 3

Proof. Recall the definitions of the events $\mathcal{E}_y$ and $\mathcal{J}$ from (38) and Section IV C, respectively. By (39) and Corollary 6 the probability that these all hold simultaneously is

$$\mathbb{P}\big[\mathcal{J} \cap \bigcap_{y\in[Y]} \mathcal{E}_y\big] \geq 1 - Ye^{-\frac{\varepsilon^2}{32}d} - 4J^2D^2e^{-\frac{\varepsilon^2}{8}d}. \tag{46}$$

If d satisfies (3) and $J, D \geq 2$, then this is

$$\geq 1 - \frac{Y}{JD} - \frac{4J^2D^2}{(4JD)^2} \geq \frac{1}{4}.$$

Fix a Π for which $\mathcal{J}$ and $\bigcap_y \mathcal{E}_y$ hold. Let $\mathcal{C}$ be the corresponding compression scheme described in Section IV A. Then by the arguments in Sections IV B and IV C, $\mathcal{C}$ satisfies the bounds in Theorem 3. □

E. Implications

1. Quantum tomography

Every empirical observation can be captured in the form of data tables $\mathcal{D}$,

$$\mathcal{D} = \begin{pmatrix} f_{1|11} & \cdots & f_{Z|11} & \cdots & f_{1|1Y} & \cdots & f_{Z|1Y} \\ f_{1|21} & \cdots & f_{Z|21} & \cdots & f_{1|2Y} & \cdots & f_{Z|2Y} \\ \vdots & & \vdots & & \vdots & & \vdots \\ f_{1|X1} & \cdots & f_{Z|X1} & \cdots & f_{1|XY} & \cdots & f_{Z|XY} \end{pmatrix}. \tag{47}$$

Here, $f_{z|xy}$ describes the empirical relative frequency for measuring the measurement outcome z given that the considered system is in a state labelled by x and given that we have performed the measurement y. If the observed system is quantum mechanical, then the states are described in terms of density matrices $(\rho_x)_{x=1}^X$ and each measurement y is described in terms of measurement operators $(E_{yz})_{z=1}^Z$. Note that the indices x, y, z are not restricted to be scalar but can be multidimensional. For instance $x = (x_1, x_2, x_3) \in \mathbb{R}^3$ if ρ_x is associated with a state preparation at spatial coordinates (x_1, x_2, x_3), or $y = t \in \mathbb{R}$ if a measurement is repeated at different times t.

Of course, in practice, $f_{z|xy}$ only approximates $\operatorname{tr}(\rho_x E_{yz})$. One way to model our uncertainty is by assuming that

$$\|\big(f_{z|xy} - \operatorname{tr}(\rho_x E_{yz})\big)_{xyz}\|_\infty \leq \delta \tag{48}$$

for $\delta > 0$, i.e., $|f_{z|xy} - \operatorname{tr}(\rho_x E_{yz})| \leq \delta$ for all (x, y, z). Ideally we would be able to characterize the set $Q[\mathcal{D}]$ of all quantum models that are compatible with $\mathcal{D}$ given the uncertainty (48). The set $Q[\mathcal{D}]$ is a subset of cartesian products of state space and measurement space. One practical characteristic of $Q[\mathcal{D}]$ is the dimension $d_{\min}$ of the lowest dimensional model in $Q[\mathcal{D}]$. For instance, if our guiding principle is Occam's razor (e.g., to counteract overfitting) then, to describe $\mathcal{D}$ theoretically, we should report a model from $Q[\mathcal{D}]$ whose dimension equals $d_{\min}$.

Denote by $\vec{\mathcal{D}}_{yz}$ the column of $\mathcal{D}$ corresponding to the measurement operator E_{yz}. The dimension $d_{\min}$ can be trivially upper bounded by X because, as already pointed out in [10], the choices ρ_x on $\mathbb{C}^X$ with

$$\rho_x = |x\rangle\langle x|$$

and E_{yz} on $\mathbb{C}^X$ with

$$(E_{yz})_{i,j} = (\vec{\mathcal{D}}_{yz})_i\, \delta_{i,j}$$

(corresponding to the factorization $\mathcal{D} = I\mathcal{D}$) is a valid model for $\mathcal{D}$, i.e.,

$$(\rho_1, ..., \rho_X, E_{11}, ..., E_{YZ}) \in Q[\mathcal{D}].$$

To see this, recall that due to normalization within $\mathcal{D}$, $\sum_z \vec{\mathcal{D}}_{yz} = (1, ..., 1)^T$ and consequently, not only are the

proposed states normalized and psd but we also have that $\sum_z E_{yz} = I$.

Finding non-trivial upper bounds on $d_{\min}$ is genuinely difficult. This is why here we address the following question: given a D-dimensional model in $Q[\mathcal{D}]$, can we—while respecting (48)—find a d-dimensional model in $Q[\mathcal{D}]$ with $d \ll D$? In other words, is it possible to *compress* the given D-dimensional model into a lower-dimensional model?

The results from sections II and IV D show that the answer is yes and no. It is affirmative for models whose measurements have few outcomes and small trace norm (see section IV D) and it is negative for measurements whose POVM elements are too long to be compressed into a compressed measurement space (see section II).

We already discussed an example for the incompressibility of quantum models in section II. Now we consider a closely related quantum model which admits exponential compression. Assume ρ_x are D-dimensional states, $Z = 2$ and assume that for all $y \in [Y]$ we have $\mathrm{tr}(E_{y1}) = 1$. This example is of practical relevance because carefully calibrated measurements are often believed to be clean in the sense that their POVM elements have trace norm $\mathcal{O}(1)$ (e.g., rank-1 projectors). In this setting, the compression $\mathcal{C}$ allows us to compress these states and measurements into a quantum model of dimension $d = \mathcal{O}(\log(D)/\varepsilon^2)$ while

$$\left|\mathrm{tr}\big(\mathcal{C}(\rho_x)\mathcal{C}(E_{y1})\big) - \mathrm{tr}(\rho_x E_{y1})\right| \leq 200\varepsilon$$

(recall (43)). It follows that

$$\left\|\Big(\mathrm{tr}(\rho_x E_{yz}) - \mathrm{tr}\big(\mathcal{C}(\rho_x)\mathcal{C}(E_{yz})\big)\Big)_{xyz}\right\|_\infty \leq 200\varepsilon \tag{49}$$

because $\mathrm{tr}(\rho_x E_{y2}) = 1 - \mathrm{tr}(\rho_x E_{y1})$. Let $\mathcal{D}$ denote the measurement data associated to the states ρ_x and the binary measurements E_{yz}.

By (48), $\big(\mathrm{tr}(\rho_x E_{yz})\big)_{xyz}$ satisfies $\|(\mathrm{tr}[\rho_x E_{yz}])_{xyz} - \mathcal{D}\|_\infty \leq \delta$. However, the precise values $\mathrm{tr}[\rho_x E_{yz}]$ are unknown to the experimentalist. Set $\varepsilon = \delta/200$. Then, for this choice for ε, the compressed states $\mathcal{C}(\rho_x)$ and measurements $\mathcal{C}(E_{xy})$ satisfy (48) and thus, they form a valid and exponentially smaller quantum model on $\mathbb{C}^d$ with $d = \mathcal{O}(\log(D)/\varepsilon^2)$.

2. *Limits of robust dimension witnessing*

A function $f : \mathcal{D} \mapsto f(\mathcal{D}) \in \mathbb{R}$ is a so-called quantum dimension witness [3] if for some $Q_D \in \mathbb{R}$,

$$f\big((\mathrm{tr}[\rho_x E_{yz}])_{xyz}\big) \leq Q_D \tag{50}$$

for all quantum states ρ_x and measurements E_{yz} whose dimension is $\leq D$. Dimension witnesses were introduced in the context of Bell inequalities [1]. In the prepare-and-measure scenario we are considering here, they have been studied extensively in the past years (see for instance [7, 10, 15, 16]). Assume ρ'_x and $E_{yz'}$ are such that

$$f\big((\mathrm{tr}[\rho'_x E'_{yz}])_{xyz}\big) > Q_k.$$

By (50), the dimension D of ρ'_x and $E_{yz'}$ satisfies $D > k$. Hence, dimension witnesses yield lower bounds on the Hilbert space dimension. Indeed, the search for dimension witnesses is motivated by the need to certify high-dimensionality of quantum systems in a device-independent manner, i.e., by only looking at the measured data $\mathcal{D}$. In the following, $f^*\big((\mathrm{tr}[\rho'_x E'_{yz}])_{xyz}\big)$ denotes the dimension lower bound that is implied by the dimension witness f.

Robustness of dimension witnesses with respect to loss has been studied in [4]. A more general objective is the analysis of the dimension witness' robustness against more general noise $\mathcal{D}_{xyz} \neq \mathrm{tr}[\rho'_x E'_{yz}]$. Here, we quantify noise in the measured data $\mathcal{D}$ in terms of l_∞-norm on $\mathbb{R}_+^{X\times YZ}$. One approach to define robustness of dimension witnesses is the demand that there exists $L > 0$ such that for all $\delta > 0$

$$|f^*\big((\mathrm{tr}[\rho'_x E'_{yz}])_{xyz}\big) - f^*\big(\mathcal{D}\big)| \leq L\|(\mathrm{tr}[\rho'_x E'_{yz}])_{xyz} - \mathcal{D}\|_\infty. \tag{51}$$

In other words, the function f^* is Lipschitz-continuous.

Recall the binary example from section IV E 1. There, the compressed model has dimension $\mathcal{O}(\log(D)/\varepsilon^2)$ even though the original model is D-dimensional. By (49) and (51),

$$\left|f^*\big((\mathrm{tr}[\rho_x E_{yz}])_{xyz}\big) - f^*\Big(\big(\mathrm{tr}[\mathcal{C}(\rho_x)\mathcal{C}(E_{yz})]\big)_{xyz}\Big)\right| \leq L\cdot 200\varepsilon \tag{52}$$

if f^* is Lipschitz-continuous. Since f^* provides a lower bound on the dimension,

$$f^*\Big(\big(\mathrm{tr}[\mathcal{C}(\rho_x)\mathcal{C}(E_{yz})]\big)_{xyz}\Big) = \mathcal{O}(\log(D)/\varepsilon^2).$$

From (52), we have

$$\begin{aligned} f^*\Big((\mathrm{tr}[\rho_x E_{yz}])_{xyz}\Big) &= \mathcal{O}(\log(D)\varepsilon^{-2} + L\varepsilon) \\ &= \mathcal{O}(\max(\log(D), L^{2/3}\log(D)^{1/3})), \end{aligned} \tag{53}$$

where in the second step we have chosen $\varepsilon = \min(1/2, L^{-1/3}(\log D)^{1/3})$.

We conclude the following for experiments whose measurements have few outcomes and small trace norm: For all Lipschitz-continuous lower bounds f^* there will be an exponential gap (53) between the dimension of the underlying Hilbert space and the dimension lower bound certified by the lower bound f^* and f respectively.

3. *Consequences for one-way quantum communication complexity*

Let $f : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$ be a Boolean function. Assume Alice gets the input $x \in \{0,1\}^n$ and Bob

gets $y \in \{0,1\}^m$. In one-way quantum communication protocols [13], Alice is allowed to send a single quantum state ρ_x to Bob. After receiving ρ_x, Bob tries to output $f(x,y)$. For that purpose, he chooses a POVM $(E_y, I - E_y)$ and measures ρ_x. If he measures E_y he sets $a = 1$. Otherwise, he sets $a = 0$. The one-way quantum bounded error communication complexity of f, denoted $Q^1(f)$, is the minimal number of qubits Alice needs to send to Bob so that $a = f(x,y)$ with probability 2/3.

To understand the connection to psd factorizations, let $A \in \mathbb{R}^{2^n \times 2^m}$ such that $A_{xy} = f(x,y)$. Hence, A is a 0/1-matrix describing f. Alice and Bob try to find the smallest d such that for every x, y there exists a state ρ_x and a POVM element E_y such that $\mathrm{tr}(\rho_x E_y) \geq 2/3$ if $A_{xy} = 1$ and $\mathrm{tr}(\rho_x E_y) \leq 1/3$ if $A_{xy} = 0$. Recalling (47), we note that this is equivalent to finding an approximate low-dimensional quantum model for the $(2^n \times 2^{m+1})$-data table

$$\mathcal{D}(A) := \big((\vec{A}_y | \vec{1} - \vec{A}_y) : y \in \{0,1\}^m\big)$$

where $\vec{A}_y$ denotes the column of A associated the input y on Bob's side. In $\mathcal{D}(A)$, the two column vectors $(\vec{A}_y | \vec{1} - \vec{A}_y)$ correspond to the binary POVM

$$(E_{y1}, E_{y2}) := (E_y, I - E_y).$$

Let P be a D-dimensional protocol to solve the one-way quantum communication task associated to a specific boolean function f. Hence, $Q^1(f) \leq \log_2(D)$. Being a protocol, P specifies mappings $x \mapsto \rho_x$ and $y \mapsto E_y$ for all inputs x, y.

We would like to apply Theorem 3 to reduce the dimension of the states and measurements used by P. This can result in significant savings whenever Bob's measurements are sufficiently unbalanced; that is, for each y, either $\mathrm{rank} E_{y1}$ or $\mathrm{rank} E_{y2}$ is small. Specifically define

$$r := \max_{y \in \{0,1\}^m} \min_{z \in \{0,1\}} \mathrm{rank} E_{yz}. \tag{54}$$

Then we will show that any bounded-error one-way protocol can be compressed to reduce the communication cost to $O(\log(nmr \log(D)))$.

Corollary 10. *Let $x \in \{0,1\}^n$ and $y \in \{0,1\}^m$. Assume that $x \mapsto \rho_x$ and $y \mapsto E_y$ is a D-dimensional protocol to solve the one-way quantum communication task associated to a Boolean function $f : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$ with error ε_0. Suppose that*

$$d \geq \frac{1280000 r^2}{\varepsilon_1^2} \max\big(\ln(4(2^n + 2^{m+1})D), r\big) \tag{55}$$

Then there exists a protocol that transmits a d-dimensional state from Alice to Bob and achieves a worst-case error of $\leq \varepsilon_0 + \varepsilon_1$.

Proof. Assume WLOG that for each y, $\mathrm{rank} E_{y1} \leq \mathrm{rank} E_{y2}$, so that $r = \max_y \mathrm{rank}(E_{y1})$. Let $\varepsilon = \varepsilon_1/200r$. Observe that $\mathrm{tr} E_{y1} \leq r$ for each y. We would like to apply Theorem 3 to this collection of measurements. Here $J = 2^n + 2^m \cdot 2$. Thus we we obtain d-dimensional compressed states $\{\rho'_x\}$ and measurements $\{E'_{yz}\}$ such that

$$\big|\mathrm{tr}(\rho'_x E'_{yz}) - \mathrm{tr}(\rho_x E_{yz})\big| \leq \varepsilon_1 \tag{56}$$

Here we also used the fact that $\mathrm{tr}(\rho'_x E'_{y2}) = 1 - \mathrm{tr}(\rho'_x E'_{y1})$, so we can extend the error bounds for the $z = 1$ measurements to $z = 2$ measurements. □

4. *Quantum message identification*

The problem of quantum message identification (introduced by Winter in [17]) can be thought of as a quantum generalization of the problem of testing whether two bit strings are equal. Alice and Bob get descriptions of D-dimensional pure states $\tau = |\phi\rangle\langle\phi|$ and $\pi = |\theta\rangle\langle\theta|$, respectively. The goal is for Bob to output a bit that equals 1 with probability close to $\mathrm{tr}(\tau\pi)$, which we think of primarily as guessing whether τ and π are approximately equal or approximately orthogonal. The problem would be trivial if Alice could send Bob the state τ, but instead she is restricted to transmitting only a d-dimensional system. (Winter considered a more general setting in which the parties can also use entanglement and/or shared or private randomness.)

Alice can use a channel $T : \mathcal{A}_1 \to \mathcal{A}_2$ to send quantum states to Bob. In quantum message identification, Bob needs to decide (up to some error) whether or not $\tau \approx \pi$. The map $\mathcal{E} : \tau \mapsto \mathcal{E}(\pi) \in \mathcal{A}_1^{\otimes n}$ describes Alice's encoding of her state π before she uses the channel T n-times to send her encoded state to Bob. Upon receiving $T^{\otimes n}\big(\mathcal{E}(\pi)\big)$, Bob applies a binary measurement $(\mathcal{D}_\tau, I - \mathcal{D}_\tau)$ to decide whether or not he should reject that $\tau \approx \pi$. Hence, $\mathcal{D}$ maps τ to a POVM element, i.e., $0 \leq \mathcal{D}_\tau \leq I$. The maps $(\mathcal{E}, \mathcal{D})$ provide an (n, λ)-quantum-ID code [17] if

$$\forall \pi, \tau \quad \big|\mathrm{tr}(\pi\tau) - \mathrm{tr}\left(T^{\otimes n}\big(\mathcal{E}(\pi)\big)\mathcal{D}_\tau\right)\big| \leq \lambda/2. \tag{57}$$

Let now $T = \mathrm{id} : \mathbb{C}^d \to \mathbb{C}^d$ be the identity channel and therefore, fix $n = 1$. Given d and λ, how large can D (i.e., the dimension of the message space) be? The following Theorem is the main result of [17].

Theorem 11 (Proposition 17 in [17])**.** *For $0 < \lambda < 1$, there exists on $\mathbb{C}^d$ a quantum-ID code of error λ such that $D = \left\lfloor d^2 \frac{(\lambda/100)^4}{4\log(100/\lambda)} \right\rfloor$.*

Equation (57) can be reinterpreted as the task to find mappings $\mathcal{E}$ and $\mathcal{D}$ to "compress" quantum states and rank-1 measurements such that the inner products between states and measurements are preserved approximately. This observation provides the link between the study of quantum message identification and the compression of quantum models. On a technical level, the

results from [17] and our results differ in two regards. Firstly, Theorem 11 does not take into account situations where both Alice and Bob get mixed quantum states as inputs. These situations are covered by our considerations. Secondly, the definition in (57) demands approximate preservation of the inner product between all possible pure states whereas we consider finite families of states and measurements. For finite sets we get exponential instead of quadratic compression. For example in case of pure quantum messages, Alice and Bob agree on a family of pure quantum messages $\mathcal{F} = \{|\psi_j\rangle\langle\psi_j|\}_j \subseteq \mathbb{C}^{D\times D}$. Then, Alice gets $\pi \in \mathcal{F}$ and Bob gets $\tau \in \mathcal{F}$. Due to the purity of π and τ, the rank constraints in (40) are not a bottleneck. Thus, we can compress into a Hilbert space with dimension $\mathcal{O}\big(\log(|\mathcal{F}|)/\varepsilon^2\big)$. The compression error is controlled by (42).

V. COMPRESSION OF A POVM WITH EXPONENTIALLY DECAYING SPECTRUM

The previous discussion assumed that with the exception of a single POVM element per measurement, all POVM elements have small rank. This constraint might be unnatural when we are asking for the compression of experimental states and measurements. It is more natural to expect that experimental measurements are full rank with thin spectral tails. The purpose of this section is to go beyond the strict low-rank assumption by assuming instead that the spectrum admits an exponentially decaying upper bound. The main result of this section is a proof of Theorem 4.

Theorem 4 (restatement). *Let ρ_x, E_{yz} and J be as in Theorem 3. Denote by $\big(\varepsilon_n^{(y)}\big)_n$ the spectrum (ordered descendingly) of $\sum_{z=1}^{Z-1} E_{yz}$. To describe potentially thin spectral tails we assume that for some $j^*, b > 0$ and for all $j > 0$,*

$$\varepsilon_{j^*+j} \leq e^{-bj}. \tag{58}$$

Let $\varepsilon \in (0, 1/2]$ and fix $d \in \mathbb{N}$ such that

$$d \geq \frac{128}{\varepsilon^2} \ln(4JD) \tag{59}$$

$$d \geq \frac{128}{\varepsilon^2}\left(j^* + \frac{1}{b}\ln\frac{8}{\varepsilon}\right) \tag{60}$$

Then, there exist d-dimensional quantum states ρ'_x and d-dimensional POVMs $\big(E'_{yz}\big)_{z=1}^{Z}$ with the same approximation promises as Theorem 3. Again, the compressed quantum model $\big(\rho'_x\big)_x$, $\big(E'_{yz}\big)_{yz}$ can be computed in randomized polynomial time in X, Y, Z and D.

Proof. We will use the same compression scheme as in previous sections. As with Theorem 3, we use Corollary 6 to argue that the ε-JL property (i.e. $\mathcal{J}$) holds with probability $\geq 1 - 4J^2D^2e^{-\frac{\varepsilon^2}{8}d}$.

We will also recall the same definition of $\mathcal{E}_y$ from (38), and attempt to prove that it holds with high probability for each y. To this end, fix a particular choice of y and set $E := I - E_{yZ}$. Let the spectral decomposition of E be

$$E = \sum_{j=1}^{\operatorname{rank}(E)} \varepsilon_j |\varepsilon_j\rangle\langle\varepsilon_j|,$$

with $\varepsilon_1 \geq \varepsilon_2 \geq \ldots$. By our assumption in (58)), after the first j^* eigenvalues, the remaining eigenvalues decay exponentially. Specifically $\varepsilon_{j^*+j} \leq e^{-bj}$.

Previously we used the low rank of the original POVM elements to control the fluctuations in the spectrum of the compressed POVM elements. Now our POVM elements will in general be full rank. Moreover, we cannot simply divide them into a low-rank piece and a low-norm piece, since in the worst case $\|\Pi E \Pi\|$ can be much large than $\|E\|$. Instead we split the exponentially decaying part of E into components each of rank

$$r := \frac{1}{b}\ln\frac{8}{\varepsilon}. \tag{61}$$

(We will see later the reason for this choice.)

$$E = \sum_{i\geq 0} E^{(i)} \tag{62}$$

$$E^{(0)} = \sum_{j\leq j^*+r} \varepsilon_j |\varepsilon_j\rangle\langle\varepsilon_j| \tag{63}$$

$$E^{(i)} = \sum_{j\in j^*+ri+[r]} \varepsilon_j |\varepsilon_j\rangle\langle\varepsilon_j| \text{ for } i > 0, \tag{64}$$

where $j^* + ri + [r]$ denotes the set $\{j^* + ri + 1, \ldots, j^* + ri + r\}$.

We will control the norm $\|\Pi E \Pi^\dagger\|$ by bounding each block separately

$$\|\Pi E \Pi^\dagger\| \leq \sum_{i\geq 0} \|\Pi E^{(i)} \Pi^\dagger\| \tag{65}$$

Now combining (60) and (61) we have

$$d \geq \frac{128}{\varepsilon^2}(j^* + r). \tag{66}$$

By Lemma 9,

$$\mathbb{P}\big[\|\Pi E^{(0)}\Pi\| \geq 1 + \varepsilon/2\big] \leq de^{-d\varepsilon^2/128}. \tag{67}$$

We will complete the proof by arguing that the remaining terms in (65) have norm $\leq \varepsilon/2$ with high probability. Using the operator inequality $E^{(i)} \leq e^{-bri}\sum_{j\in j^*+ri+[r]} |\varepsilon_j\rangle\langle\varepsilon_j|$, we have

$$\left\|\sum_{i\geq 1} E^{(i)}\right\| \leq \sum_{i\geq 1} e^{-bri} \left\|\Pi \sum_{j\in j^*+ri+[r]} |\varepsilon_j\rangle\langle\varepsilon_j| \Pi^\dagger\right\| . \tag{68}$$
$$=: \sum_{i\geq 1} e^{-bri} \left\|G_i G_i^\dagger\right\| =: X$$

In the second line, we have defined complex $d \times r$-dimensional matrices G_i which are i.i.d. and each comprised of i.i.d. complex Gaussians with variance $1/d$, i.e. as in Theorem 8. Define

$$\bar{X} = \frac{e^{-br}}{1-e^{-br}}(1+\sqrt{r/d})^2 \tag{69}$$

Now we bound the moment-generating function of X (assuming $t \leq d/2$) by

$$\begin{aligned}
\mathbb{E}[e^{tX}] &= \prod_{i\geq 1} \mathbb{E}[\exp(te^{-bri}\|G_i G_i^\dagger\|] \\
&\leq \prod_{i\geq 1} \mathbb{E}[\exp(t\|G_i G_i^\dagger\|]^{e^{-bri}} \\
&\leq \prod_{i\geq 1} \left[d\exp(t(1+\sqrt{r/d})^2 + \frac{t^2}{d}(1+r/d))\right]^{e^{-bri}} \\
&= \left[d\exp(t(1+\sqrt{r/d})^2 + \frac{t^2}{d}(1+r/d))\right]^{\frac{e^{-br}}{1-e^{-br}}} \\
&= e^{t\bar{X}} d^{\frac{e^{-br}}{1-e^{-br}}} \exp\left(\frac{t^2}{d}(1+r/d)\frac{e^{-br}}{1-e^{-br}}\right)
\end{aligned}$$

We have used here first the independence of each G_i, then the bound $\mathbb{E}[x^\alpha] \leq \mathbb{E}[x]^\alpha$ for $x \geq 0, 0 \leq \alpha \leq 1$, then Theorem 8 and the definition of $\bar{X}$. From (61) we have $e^{-br} = \varepsilon/8$ and thus $\frac{e^{-br}}{1-e^{-br}} \leq \varepsilon/4$. Using also (66) we have $(1+r/d)\frac{e^{-br}}{1-e^{-br}} \leq \varepsilon/4$ and $\bar{X} \leq \varepsilon/4$. Now we choose $t = d/2$ and use Markov's inequality to bound

$$\begin{aligned}
\mathbb{P}\big[X \geq \bar{X} + \frac{\varepsilon}{4}\big] &\leq \mathbb{E}\left[\exp(tX)\right] e^{-t(\bar{X}+\frac{\varepsilon}{4})} \\
&\leq d^{\varepsilon/4} \exp\left(t\bar{X} + \frac{t^2}{d}\frac{\varepsilon}{4}\right) e^{-t(\bar{X}+\frac{\varepsilon}{4})} \\
&\leq d^{\varepsilon/4}\exp[-d\varepsilon/16] \\
&= \exp[-d\varepsilon/20].
\end{aligned}$$

Since $\bar{X} \leq \varepsilon/4$ we conclude that

$$\left\|\sum_{i\geq 1} E^{(i)}\right\| \leq \varepsilon/2, \tag{70}$$

with probability $\geq 1 - \exp[-d\varepsilon/20] \geq 1 - d\exp[-d\varepsilon^2/128]$. By the union bound, both (70) and $\|\Pi E^{(0)}\Pi^\dagger\| \leq 1 + \varepsilon/2$ hold with probability

$$\geq 1 - 2d\exp[-d\varepsilon^2/128] \geq 1 - \frac{2d}{4JD}. \tag{71}$$

This lower bounds the probability of a single $\mathcal{E}_y$ holding.

By the union bound, we have

$$\mathbb{P}\big[\mathcal{J} \cup \bigcup_{y\in Y} \mathcal{E}_y\big] \geq 1 - 4J^2D^2 e^{-d\frac{\varepsilon^2}{8}} - \frac{2dY}{4JD} > 0. \tag{72}$$

The rest of the proof is the same as in Theorem 3. □

VI. CONCLUSION

Nonnegative matrices $\mathcal{D} \in \mathbb{R}^{X\times YZ}$ can be described in terms of positive semidefinite factorizations or, in some cases, in terms of quantum models. The number of degrees of freedom in both of these models is a function of their dimension D. In practice it is often not necessary to describe $\mathcal{D}$ in terms of an exact psd or quantum factorization but it is sufficient to describe some matrix $\mathcal{D}'$ with $\|\mathcal{D} - \mathcal{D}'\|_\infty \leq \delta$. It is then natural to ask for the minimal dimension of such approximate factorizations of $\mathcal{D}$.

Assume A_i and B_j constitute a D-dimensional psd factorization of $\mathcal{D}$. In this paper we showed that in randomized polynomial time we can find d-dimensional psd matrices A'_i and B'_j such that

$$\begin{aligned}
\big|\mathcal{D}_{i,j} - \operatorname{tr}(A'_i B'_j)\big| &\leq 192\varepsilon \operatorname{tr}(A_i)\operatorname{tr}(B_j) \\
d &= \left\lceil \frac{32}{\varepsilon^2}\ln\big(2JD\big)\right\rceil
\end{aligned} \tag{73}$$

where $J = X + YZ$ (cf. Theorem 2). Hence, psd factorizations with bounded trace norm admit exponential compression.

Quantum models are psd factorizations with additional normalization constraints. These normalization constraints can affect compressibility dramatically. Consider for instance $\mathcal{D} = I \in \mathbb{R}^{D\times D}$ with the property $\mathcal{D}_{ij} = \operatorname{tr}\big(|i\rangle\langle i||j\rangle\langle j|\big)$. This factorization is both psd and quantum ($Y = 1$; $E_{1z} = |z\rangle\langle z|$). As a psd factorization, it admits exponential compression. However, regarded as a quantum factorization, every δ-approximate, d-dimensional quantum factorization of $\mathcal{D}$ satisfies

$$d \geq D - Z\delta = (1-\delta)D.$$

This is an application of the lower bound derived in section II (cf. Theorem 1). Aiming for the compression of quantum models, this observation might seem discouraging.

However, this lower bound from Theorem 1 only prevents compression below Z. It thus remains to investigate compressibility of quantum models with $Z = \mathcal{O}(1)$ in D. Introducing a compression scheme which runs in randomized polynomial time in X, Y and Z, we showed (cf. Theorem 3) that there exist d-dimensional states ρ'_x and POVMs E'_y with approximation promises analogous to (73) and

$$d = \max\Big\{\frac{16}{\varepsilon^2}\ln(4JD), \frac{32}{\varepsilon^2}\operatorname{rank}\Big(\sum_{z=1}^{Z-1} E_{yz}\Big)\Big\}. \tag{74}$$

Note that this compression scheme thus almost achieves the lower bound if $\operatorname{rank}\big(\sum_{z=1}^{Z-1} E_{yz}\big) \sim Z$.

Equation (74) is only meaningful if the relevant measurements are low rank. Experimental measurements are however expected to only be approximately low rank. To cover these scenarios we derived a version of Theorem 3

which can handle measurements that are only approximately low rank in the sense that their spectrum decays exponentially (cf. Theorem 4).

In this paper we only started the exploration of applications of these results. We briefly commented on their implications in the fields quantum tomography, dimension witnessing, one-way quantum communication complexity and quantum message identification. In particular our results imply that

- high-dimensional models whose measurements are low rank and
- low-dimensional models

are equivalent from an operational perspective if the number of outcomes per measurement is small.

On a less technical and more philosophical side, our results open up a path towards demystification of the apparent miracle that our inherently complex, high-dimensional world sometimes admits simple, low-dimensional descriptions which form the basis of science. Since our compression schemes are essentially just random projections, these low-dimensional descriptions can potentially be regarded as random projections. If that is the case then it should not be surprising that simple, effective, low-dimensional models can be found if the measurements data stems from few-outcome and sharp/clean measurements, i.e., measurements which are approximately low rank.

ACKNOWLEDGMENT

We would like to thank Guillaume Aubrun, Matt Coudron, David Gross, Kristan Temme and Henry Yuen for interesting discussions and for being such amazing colleagues. CS acknowledges support through the SNSF postdoctoral fellowship, support from the SNSF through the National Centre of Competence in Research "Quantum Science and Technology" and funding by the ARO grant Contract Number W911NF-12-0486. AWH was funded by NSF grant CCF-1111382 and ARO contract W911NF-12-1-0486.

[1] A. Acín, N. Brunner, N. Gisin, S. Massar, S. Pironio, and V. Scarani. Device-independent security of quantum cryptography against collective attacks. *Physical Review Letters*, 98(23):230501, 2007.

[2] A. Barvinok. *A course in convexity*, volume 54. American Mathematical Soc., 2002.

[3] N. Brunner, M. Navascués, and T. Vértesi. Dimension witnesses and quantum state discrimination. *Physical Review Letters*, 110(15):150501, 2013.

[4] M. Dall Arno, E. Passaro, R. Gallego, and A. Acín. Robustness of device-independent dimension witnesses. *Physical Review A*, 86(4):042312, 2012.

[5] H. Fawzi, J. Gouveia, P. A. Parrilo, R. Z. Robinson, and R. R. Thomas. Positive semidefinite rank, 2014, arXiv:1407.4095.

[6] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. de Wolf. Linear vs. semidefinite extended formulations: exponential separation and strong lower bounds. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 95–106. ACM, 2012.

[7] R. Gallego, N. Brunner, C. Hadley, and A. Acín. Device-independent tests of classical and quantum dimensions. *Physical review letters*, 105(23):230501, 2010.

[8] J. Gouveia, P. A. Parrilo, and R. R. Thomas. Approximate cone factorizations and lifts of polytopes, 2013, arXiv:1308.2162.

[9] U. Haagerup and S. Thorbjørnsen. Random matrices with complex gaussian entries. *Expositiones Mathematicae*, 21(4):293 – 337, 2003.

[10] N. Harrigan, T. Rudolph, and S. Aaronson. Representing probabilistic data via ontological models, 2007, arXiv:0709.1149.

[11] A. W. Harrow, A. Montanaro, and A. J. Short. Limitations on quantum dimensionality reduction. In *Automata, Languages and Programming*, pages 86–97. Springer, 2011.

[12] W. B. Johnson and J. Lindenstrauss. Extensions of lipschitz maps into a hilbert space. *Contemporary Mathematics*, 26:189–206, 1984.

[13] H. Klauck. Quantum communication complexity. In *ICALP*, pages 241–252, 2000, arXiv:quant-ph/0005032.

[14] T. Lee, Z. Wei, and R. de Wolf. Some upper and lower bounds on psd-rank. *arXiv preprint arXiv:1407.4308*, 2014.

[15] C. Stark. Self-consistent tomography of the state-measurement gram matrix. *Physical Review A*, 89(5):052109, 2014.

[16] S. Wehner, M. Christandl, and A. C. Doherty. Lower bound on the dimension of a quantum system given measured data. *Physical Review A*, 78(6):062112, 2008.

[17] A. Winter. Quantum and classical message identification via quantum channels. *arXiv preprint quant-ph/0401060*, 2004.